



10-19-00

A

THE ASSISTANT COMMISSIONER OF PATENTS
Washington, D.C. 20231

DOCKET NUMBER: RPS920000076US1
October 18, 2000

Sir:

Transmitted herewith for filing is the Patent Application of:

Inventors: Charles David Bauman, Richard Bealkowski, Thomas J. Clement, Jerry William Pearce, Michael Robert Turner

For: SUPPORT FOR EXHAUSTION RECOVERY IN A DATA PROCESSING SYSTEM WITH MEMORY MIRRORING

Enclosed are:

jc893 U.S. PTO
09/691533
10/18/00

Patent Specification and Declaration
 4 sheets of drawing(s).
 An assignment of the invention to International Business Machines Corporation (includes Recordation Form Cover Sheet).
 A certified copy of a _____ application.
 Information Disclosure Statement, PTO 1449 and copies of references.

The filing fee has been calculated as shown below:

| For | Number Filed | Number Extra | Rate | Fee |
|--|-----------------|-----------------|----------|----------|
| Basic Fee | | | | \$710.00 |
| Total Claims | 21 | - 20 | 1 x 18 = | \$ 18.00 |
| Indep. Claims | 5 | - 3 | 2 x 80 = | \$160.00 |
| MULTIPLE DEPENDENT CLAIM PRESENTED x 260 = | | | | \$ |
| | | | TOTAL | \$888.00 |

Please charge IBM Corporation Deposit Account No. 50-0563 in the amount of \$888.00. A duplicate copy of this sheet is enclosed.

The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to IBM Corporation Deposit Account 50-0563. A duplicate copy of this sheet is enclosed.

Any additional filing fees required under 37 CFR §1.16.

Any patent application processing fees under 37 CFR §1.17.

| | |
|--|--|
| CERTIFICATE OF MAILING BY "EXPRESS MAIL" UNDER 37 CFR § 1.10 | |
| "Express Mail" mailing label number EL 453462924US | |
| Date of Mailing October 18, 2000 | |
| I hereby certify that the documents indicated below are being deposited with the United States Postal Service under 37 CFR 1.10 on the date indicated above and are addressed to Box Patent Applications, Assistant Commissioner of Patents, Washington, D.C. 20231 and mailed on the above Date of Mailing with the above "Express Mail" mailing label number | |
| Dena Faris (name of person mailing paper) | SIGNATURE of person mailing paper or fee |

Respectfully submitted,

By:
Michael R. Barré
Registration No. 44,023
FELSMAN, BRADLEY, VADEN,
GUNTER & DILLON, LLP
Suite 350 Lakewood on the Park
7600B North Capital of Texas Highway
Austin, Texas 78731
Telephone (512) 343-6116

**SUPPORT FOR EXHAUSTION RECOVERY
IN A DATA PROCESSING SYSTEM WITH MEMORY MIRRORING**

BACKGROUND OF THE INVENTION

5

1. Technical Field:

The present invention relates in general to data processing systems and in particular to data processing systems that support memory mirroring. Still more particularly, the present invention relates to systems for managing memory in data processing systems that support memory mirroring.

15 2. Description of the Related Art:

In the field of data processing systems, a design feature that can be employed in high-availability data processing systems is memory mirroring. Memory mirroring maintains a redundant copy of system memory. For example, if a system with two banks of random access memory (RAM) has memory mirroring, the first bank (e.g., Bank A) might be utilized as the primary bank, and the second bank (e.g., Bank B) might be utilized as a backup or mirror bank. The mirror bank serves as a repository for an extra copy of the data contained in the primary bank. In the event of errors in primary Bank A, the system avoids an interruption in service by utilizing mirror Bank B in lieu of primary Bank A. An error message may then be generated indicating that Bank A requires service.

A disadvantage associated with memory mirroring, however, is that less of the installed memory is available

for general use. For example, in the above example, Bank A is available for general use, but Bank B is not. When all of the primary memory is mirrored onto mirror memory, the system is said to have full mirroring, in which case only half of the installed memory is available for general use. 5 Alternatively, partial mirroring can be implemented. For example, a system with four banks of memory, (e.g., Banks A through D) may be configured to mirror Bank A onto Bank B but not to mirror Banks C and D. Partial mirroring 10 provides for general use of a larger percentage of installed memory but less protection against memory errors.

Another design feature that can be employed to provide mirrored systems (as well as non-mirrored systems) with more apparent memory for general use is memory compression. A typical system with memory compression has a compression engine that utilizes the system's installed memory (i.e., physical memory) to support a larger amount of apparent memory (i.e., real memory). For example, if a 15 system includes 32 MB of physical memory and a compression engine configured to provide a compression ratio of 2:1, the compression engine will report real memory of 64 MB. Consequently, if a system with 32 MB of installed memory has full mirroring and memory compression at a ratio of 20 2:1, the effects of the compression and mirroring offset each other. That is, the system will have a real memory 25 size that equals its physical memory size.

A danger associated with memory compression, 30 however, is that the configured compression ratio may not be obtainable for certain types of data. For example, if the data has already been compressed by software (e.g., if the data is a Joint Photographic Experts Group (JPEG) file), the compression engine may not be able to further

compress the file enough to provide a 2:1 ratio. Furthermore, if the compression engine fails to achieve the desired compression ratio and substantially all of the real memory is needed, the system may be unable to provide the required amount of real memory. When a system lacks sufficient physical memory to support all of the required real memory, the system is said to have suffered a condition known as memory exhaustion. However, memory exhaustion may also occur in systems without compression, for example if the system simply has insufficient physical memory to perform required processing.

In conventional systems, exhausting physical memory is a critical system problem. That problem is typically signaled via an appropriately high priority event, such as a non-maskable interrupt (NMI) or a form of machine check. Consequently, in conventional systems, when memory compression and memory mirroring are used together, there will generally be a decreased risk of system failure due to memory errors but an increased risk of system failure due to memory exhaustion.

Consequently, what is needed is a way to reduce the risk of system failure due to memory exhaustion, especially in systems that utilize memory compression.

SUMMARY OF THE INVENTION

According to the present invention, a memory exhaustion condition is handled in a data processing system having first and second regions of physical memory. In one particular method according to the present invention, a memory exhaustion condition is detected while the second region is mirroring at least part of the first region. In response to the memory exhaustion condition, memory mirroring is at least partially deactivated. The memory exhaustion condition is then eliminated by augmenting the first region with at least part of the second region.

In an illustrative embodiment, the data processing system compresses real memory into the first region of physical memory, and the memory exhaustion condition is detected when the first region lacks sufficient available capacity to accommodate current requirements for real memory. The memory exhaustion condition is eliminated by compressing at least part of the real memory into the second region.

All objects, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use and further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figures 1A and **1B** depict block diagrams explaining the relationship between physical memory and real memory with mirroring enabled (in **Figure 1A**) and disabled (in **Figure 1B**) according to an illustrative embodiment of the present invention;

Figure 2 illustrates a block diagram of an exemplary data processing system with facilities, in accordance with the present invention, for handling memory exhaustion conditions;

Figure 3 is a table illustrating possible settings of memory-bank select bits for configuring mirroring according to an illustrative embodiment of the present invention;

Figure 4 is a table depicting possible settings of a mirror-mode read bit for configuring mirroring according to an illustrative embodiment of the present invention; and

Figure 5 illustrates a high-level logical flowchart of an exemplary process, according to the present invention, for handling memory exhaustion conditions.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENT

With reference now to the figures, and in particular with reference to **Figures 1A** and **1B**, there are 5 depicted two block diagrams explaining the relationship between physical memory and real memory in a data processing system with a memory subsystem that supports mirroring in accordance with a relatively uncomplicated illustrative embodiment of the present invention.

10 Specifically, the diagrams relate to a data processing system that has two banks of memory of equal capacity, Bank A **10** and Bank B **12**. In **Figure 1A**, the data processing system is configured in such a way that Bank A **10** serves as a primary bank and Bank B **12** serves as a mirror of Bank A 15 **10**. In addition, the data processing system supports memory compression, such that the memory subsystem reports an address space for real memory **14** that exceeds the address space of the primary bank (i.e., Bank A **10**).

20 To describe the relationship between primary memory and real memory more precisely, this document utilizes the following symbols in the following manner:

X is the size of the physical, primary memory;

25 X(m) is the size of the physical memory serving as a mirror;

R is the compression ratio that the memory subsystem seeks to achieve (expressed as real-memory:primary-memory); and

30 Rx is the size of the real memory, as reported by the memory subsystem.

In **Figure 1A**, R equals 2:1 and X equals X_1 ; consequently, Rx equals $2 * X_1$. Also, X(m) equals X_2 .

Referring now to **Figure 1B**, there is depicted a block diagram relating to the data processing system of **Figure 1A** with, however, the memory subsystem having been reconfigured so that Bank B **12** serves as primary memory, not mirror memory. According to the present invention, such a reconfiguration may be performed automatically upon the occurrence of a memory exhaustion condition, as described in greater detail below. In the present illustrative embodiment, as a result of the reconfiguration, R equals 1:1 and X equals $X_1 + X_2$; consequently, R_X equals $1 * (X_1 + X_2)$. Also, $X(m)$ equals 0. Thus, additional physical memory is allocated to support needs for real memory, at the expense of loss of mirroring.

With reference now to **Figure 2**, there is illustrated a block diagram of a slightly more complex exemplary data processing system **20** with facilities for handling memory exhaustion conditions in accordance with the present invention. As illustrated, data processing system **20** includes a system processor **22** that communicates with a memory subsystem **24** via a host bus **26**. Memory subsystem **24** includes a cache and memory controller (CMC) **30** which supports an L3 cache **32**. In addition, CMC **30** communicates with four banks of installed random access memory (RAM), Bank 0 **40**, Bank 1 **41**, Bank 2 **42**, and Bank 3 **43**, both directly and via a mirror mode controller (MMC) **50**. Other components (not illustrated) of data processing system **20** may communicate with memory subsystem **24** via a peripheral component interconnect (PCI) bus **52**.

MMC **50** contains a three-bit latch that can be written and read by system processor **22** to alter the

mirroring characteristics of memory subsystem **24**. The latch appears as a normal I/O port.

Referring now also to **Figure 3**, two bits of that latch are bank-select bits. By setting or clearing the bank-select bits to achieve the bit patterns illustrated in **Figure 3**, system processor **22** can alter how Bank 3 **43** is used. Bank 3 **43** is therefore said to have variable addressing. In particular, when at least one of bank-select bits 0 and 1 is cleared, Bank 3 **43** serves as a mirror bank for one of the primary banks. Such a configuration therefore implements partial-mirroring. By contrast, when bank-select bits 0 and 1 are both set, mirroring is deactivated, and Bank 3 **43** serves instead as another primary bank, for example as the high bank in a four bank array.

Referring now also to **Figure 4**, the third bit in the latch is a mirror-read bit which, when clear, causes memory reads to utilize a primary bank and, when set, causes memory reads to utilize the mirror bank in lieu of the corresponding primary bank. At this point, it may be noted that, if Bank 1 **41** and Bank 2 **42** were null and the capacities of Bank 0 **40** and Bank 3 **43** were equal, the system depicted in **Figure 2** would resemble or match the system described with reference to **Figures 1A** and **1B**.

The operation of memory subsystem **24** is governed by a memory manager **60** which, in the illustrative embodiment, is implemented as part of the operating system. Accordingly, memory manager **60** is loaded into RAM during the system boot process, along with other operating system components.

With reference now to **Figure 5**, there is illustrated a high-level logical flowchart of an exemplary process, according to the present invention, for handling memory exhaustion conditions within data processing system **20**. The process begins at block **70** with data processing system **20** being powered up. Then, as part of the boot process, memory subsystem **24** is initialized, as depicted at block **72**. As part of the initialization process, CMC **30** determines how much physical memory is installed and memory manager **60** configures the bank-select and mirror-read bits within MMC **50** according to a default configuration. In the current exemplary process, MMC **50** will be configured for partial mirroring with Bank 0 **40** mirrored onto Bank 3 **43**.

A compression ratio may also be determined during this initialization process. In one embodiment, that ratio is determined based on historical compression performance, so that the initialized compression ratio may be automatically adjusted over time to more closely match the compression ratios that have actually been achieved during one or more prior operating cycles. In the present example, the compression ratio will be set at 2:1.

After the memory subsystem has been initialized, data processing system **20** begins normal operation, as depicted at block **74**, with the memory subsystem operating as configured in the initialization process. During normal operation, combinational logic controls the generation of the chip select and output enable signals for each memory bank. This logic utilizes the high-order memory address lines, the bank-select bits, the system R/W line, and the mirror-read bit to generate the required signals and timing. In mirror mode, all writes to a mirrored primary

bank (here, Bank 0 **40**) also write to the corresponding mirror bank (here, Bank 3 **43**). In addition, CMC **30** maintains usage tables for tasks such as allocating unused memory as needed and properly mapping real memory to physical memory.

5

10

15

20

25

30

During processing, CMC **30** utilizes L3 cache **32** to obtain recently accessed data, as data can be read from L3 cache **32** more quickly than from RAM. However, when the required data is not in L3 cache **32**, (i.e., when there is a cache miss), it is necessary to access RAM. Specifically, in the present illustrative embodiment, cache misses cause CMC **30** to read a compressed block of data from RAM, decompress the data, and load the data into the L3 cache. For a cache flush operation, the data to be written is compressed and written to RAM. The memory usage tables of CMC **30** are updated accordingly.

As illustrated at block **76**, while system processor **22** executes instructions, memory manager **60** monitors memory subsystem **24** to detect whether a memory exhaustion condition exists (i.e., whether the primary memory lacks sufficient available capacity to accommodate current requirements for real memory). For example, CMC **30** may be configured to generate a non-maskable interrupt or machine check when there simply is not enough physical memory to store the compressed contents of real memory, due to, for example, a request for additional real memory or a modification of existing memory contents. Alternatively, it may be determined that a memory exhaustion condition exists when current requirements for real memory cause the remaining capacity of unused physical memory to fall below a certain threshold.

If no memory exhaustion condition is detected, processing continues with mirroring, as indicated by the arrow returning to block **74**. However, if a memory exhaustion condition is detected, the process passes to block **80**, which depicts memory manager **60** recording information pertaining to the memory exhaustion condition for use in adapting the configured compression ratio to align more closely with compression ratios that have actually been achieved. Specifically, in the current exemplary process, memory manager **60** computes a new compression ratio and stores that compression ratio in non-volatile storage. Acceptable storage mediums for this data include, without limitation, non-volatile (e.g., battery-backed) RAM (i.e., NVRAM), complementary metal-oxide semiconductor (CMOS) RAM, flash memory, and magnetic or optical disks. In alternative embodiments, however, other types of exhaustion information may be stored.

After a memory exhaustion condition has been detected, memory manager **60** also reconfigures CMC **30** and MMC **50** (including setting both bank-select bits) to deactivate mirroring and convert Bank 3 **43** from a mirror bank to a primary bank, preferably as the high bank of a four-bank array, as depicted at blocks **82** and **84**. As a result, the amount of available primary memory is expanded, which eliminates the memory exhaustion condition. Then, as illustrated at block **86**, system processor **22** resumes normal operation, utilizing the expanded primary memory.

Also, as shown at block **88**, after normal processing has resumed, memory manager **60** monitors conditions to determine whether the expanded primary memory is still needed. If so, processing continues, as indicated

by the arrow leading back to block **86**. However, if real memory requirements have subsided to the point that Bank 3 **43** is no longer needed as a primary bank, memory manager **60** reconfigures CMC **30** and MMC **50** to reactivate mirroring, as depicted at block **90**. Within the process of reactivating mirroring, the mirror bank is repopulated with data to match the corresponding primary bank. This repopulation can be performed by software or by hardware. System processor **22** then resumes normal operations, as indicated by the arrow leading back to block **74**.

As has been described, the present invention handles memory exhaustion conditions by dynamically reconfiguring mirror banks to be primary banks. In addition, mirroring is reactivated when memory needs subside sufficiently. Furthermore, in data processing systems that provide memory compression, the invention provides for the automatic adjustment of the startup compression ratio, based on how much compression has actually been realized in a particular data processing system.

While the invention has been particularly shown and described with reference to one or more illustrative embodiments, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention. For example, particular illustrative embodiments involving full and partial mirroring have been present. However alternative embodiments with different numbers of memory banks and different mirroring proportions are supported as well, including data processing systems with N-way mirroring (i.e., data processing system with two

or more mirror copies of data from one or more primary memory banks). Accordingly, the present invention also supports partial deactivation of mirroring to eliminate memory exhaustion conditions. For example, if a system is initially configured to utilize two or more banks of memory for mirroring, the system may respond to memory exhaustion conditions by reconfiguring one mirror bank as a primary bank and leaving mirroring active for the other mirror bank. Furthermore, although the illustrative embodiments provide memory compression, the present invention can also be utilized to advantage in systems without compression.

Also, although certain architectures have been utilized to illustrate exemplary embodiments, the present invention could be utilized with other architectures, as well. For example, the present invention could be implemented in data processing systems including, without limitation, general-purpose personal computers, multi-processor systems, mini-computers, and mainframe computers.

In addition, various operations of the present invention have been described as being performed by various components, including memory manager **60**, CMC **30**, and MMC **50**. However, those of ordinary skill in the art will appreciate that, to a certain extent, the operations performed according to the present invention could be distributed among those components differently. In addition, the implementation of certain components as hardware, software, and/or firmware can be altered without departing from the spirit of the present invention.

For example, memory manager **60** could be implemented partially or entirely in hardware and/or

firmware. Likewise, rather than having a memory controller that generates signals such as interrupts to be processed by software, a data processing system with hardwired logic that controls mirroring may include hardwired detection logic that detects memory exhaustion conditions, as well as configuration logic that is hardwired to the detection logic in such a way that the configuration logic automatically disables mirroring and converts mirror banks into primary banks when necessary.

10

Also, although aspects of the present invention have been described with respect to a computer system executing software that directs the functions of the present invention, it should be understood that the present invention may alternatively be implemented as a program product for use with a data processing system. Programs defining the functions of the present invention can be delivered to a data processing system via a variety of signal-bearing media, which include, without limitation, non-rewritable storage media (e.g., CD-ROM), rewritable storage media (e.g., a floppy diskette or hard disk drive), and communication media, such as digital and analog networks. It should be understood, therefore, that such signal-bearing media, when carrying or encoding computer readable instructions that direct the functions of the present invention, represent alternative embodiments of the present invention.

15
20
25

Claims:

What is claimed is:

1. 1. A method of handling a memory exhaustion condition in
2 a data processing system having first and second regions of
3 physical memory, said method comprising:

4 detecting a memory exhaustion condition while
5 said second region is mirroring at least part of said first
6 region;

7 in response to said memory exhaustion condition,
8 at least partially deactivating memory mirroring between
9 said first and second regions; and

10 augmenting said first region with at least part
11 of said second region, such that said memory exhaustion
12 condition is eliminated.

1. 2. The method of Claim 1, said data processing system
2 compressing real memory into said first region of physical
3 memory, wherein:

4 said step of detecting a memory exhaustion
5 condition comprises determining that said first region
6 lacks sufficient available capacity to accommodate current
7 requirements for real memory; and

8 said step of augmenting said first region
9 comprises compressing at least part of said required real
10 memory into said at least part of said second region.

1 3. The method of Claim 2, further comprising:

2 in response to said memory exhaustion condition,
3 recording data relating to said memory exhaustion
4 condition; and

5 subsequently utilizing said data to set a
6 compression ratio for memory compression.

1 4. The method of Claim 1, said data processing system
2 having a mirror mode controller that directs a memory
3 access to said first region or said second region in
4 accordance with one or more mirror mode bits, wherein:

5 said step of augmenting said first region
6 comprises modifying at least one of said one or more mirror
7 mode bits.

1 5. The method of Claim 1, further comprising:

2 after memory mirroring has been at least
3 partially deactivated, determining that memory needs have
4 diminished; and

5 in response, at least partially reactivating
6 mirroring.

1 6. A data processing system comprising:
2 first and second regions of physical memory;
3 detection logic that detects a memory exhaustion
4 condition while said second region is mirroring at least
5 part of said first region; and
6 configuration logic that, responsive to said
7 memory exhaustion condition, at least partially deactivates
8 memory mirroring between said first and second regions and
9 augments said first region with at least part of said
10 second region, such that said memory exhaustion condition
11 is eliminated.

1 7. The data processing system of Claim 6, wherein:
2 said data processing system compresses real
3 memory into said first region of physical memory;
4 said detection logic comprises a memory
5 controller that detects said memory exhaustion condition by
6 determining that said first region lacks sufficient
7 available capacity to accommodate current requirements for
8 real memory; and
9 said configuration logic comprises a memory
10 manager that augments said first region by configuring said
11 memory controller to compress at least part of said real
12 memory into said at least part of said second region.

1 8. The data processing system of Claim 7, wherein:
2 in response to said memory exhaustion condition,
3 said memory manager records data relating to said memory
4 exhaustion condition; and
5 said memory manager subsequently utilizes said
6 data to set a compression ratio for memory compression.

1 9. The data processing system of Claim 6, wherein:

2 said data processing system further comprises a
3 mirror mode controller that directs a memory access to said
4 first region or said second region in accordance with one
5 or more mirror mode bits; and

6 said configuration logic augments said first
7 region by modifying at least one of said one or more mirror
8 mode bits.

1 10. The data processing system of Claim 6, wherein:

2 after memory mirroring has been at least
3 partially deactivated, said configuration logic determines
4 that memory needs have diminished; and

5 in response, said configuration logic at least
6 partially reactivates mirroring.

1 11. A program product that handles a memory exhaustion
2 condition in a data processing system having a memory
3 controller and first and second regions of physical memory,
4 said program product comprising:

5 a computer usable medium encoding a memory
6 manager;

7 deactivation instructions within said memory
8 manager that at least partially deactivate memory mirroring
9 between said first and second regions in response to a
10 memory exhaustion condition arising while said second
11 region is mirroring at least part of said first region; and

12 augmentation instructions within said memory
13 manager that cause said memory controller to augment said
14 first region with at least part of said second region, such
15 that said memory exhaustion condition is eliminated.

1 12. The program product of Claim 11, said data processing
2 system compressing real memory into said first region of
3 physical memory, and said memory controller detecting said
4 memory exhaustion condition by determining that said first
5 region lacks sufficient available capacity to accommodate
6 current requirements for real memory, wherein:

7 said memory manager augments said first region by
8 configuring said memory controller to compress at least
9 part of said real memory into said at least part of said
10 second region.

1 13. The program product of Claim 12, wherein:

2 in response to said memory exhaustion condition,
3 said memory manager records data relating to said memory
4 exhaustion condition; and

5 said memory manager subsequently utilizes said
6 data to set a compression ratio for memory compression.

1 14. The program product of Claim 11, wherein:

2 said data processing system further comprises a
3 mirror mode controller that directs a memory access to said
4 first region or said second region in accordance with one
5 or more mirror mode bits; and

6 said memory manager augments said first region by
7 modifying at least one of said one or more mirror mode
8 bits.

1 15. The program product of Claim 11, wherein:

2 after memory mirroring has been at least
3 partially deactivated, said memory manager determines that
4 memory needs have diminished; and

5 in response, said memory manager at least
6 partially reactivates mirroring.

1 16. A memory management system that handles a memory
2 exhaustion condition in a data processing system having
3 first and second regions of physical memory, said memory
4 management system comprising:

5 detection logic that detects a memory exhaustion
6 condition while said second region is mirroring at least
7 part of said first region; and

8 configuration logic that, responsive to said
9 memory exhaustion condition, at least partially deactivates
10 memory mirroring between said first and second regions and
11 augments said first region with at least part of said
12 second region, such that said memory exhaustion condition
13 is eliminated.

1 17. The memory management system of Claim 16, said data
2 processing system compressing real memory into said first
3 region of physical memory, wherein:

4 said detection logic comprises a memory
5 controller that detects said memory exhaustion condition by
6 determining that said first region lacks sufficient
7 available capacity to accommodate current requirements for
8 real memory; and

9 said configuration logic comprises a memory
10 manager that augments said first region by configuring said
11 memory controller to compress at least part of said real
12 memory into said at least part of said second region.

1 18. The memory management system of Claim 17, wherein:
2 in response to said memory exhaustion condition,
3 said memory manager records data relating to said memory
4 exhaustion condition; and
5 said memory manager subsequently utilizes said
6 data to set a compression ratio for memory compression.

1 19. The memory management system of Claim 16, wherein:
2 said memory management system further comprises a
3 mirror mode controller that directs a memory access to said
4 first region or said second region in accordance with one
5 or more mirror mode bits; and
6 said configuration logic augments said first
7 region by modifying at least one of said one or more mirror
8 mode bits.

1 20. The memory management system of Claim 16, wherein:
2 after memory mirroring has been at least
3 partially deactivated, said configuration logic determines
4 that memory needs have diminished; and
5 in response, said configuration logic at least
6 partially reactivates mirroring.

1 21. A data processing system comprising:

2 first and second regions of physical memory;

3 detection means for detecting a memory exhaustion
4 condition while said second region is mirroring at least
5 part of said first region; and

6 configuration means, responsive to said memory
7 exhaustion condition, for at least partially deactivating
8 memory mirroring between said first and second regions and
9 augmenting said first region with at least part of said
10 second region, such that said memory exhaustion condition
11 is eliminated.

 2000 3000 4000 5000 6000 7000 8000 9000 10000 11000 12000 13000 14000 15000 16000 17000 18000 19000 20000 21000 22000 23000 24000 25000 26000 27000 28000 29000 30000 31000 32000 33000 34000 35000 36000 37000 38000 39000 40000 41000 42000 43000 44000 45000 46000 47000 48000 49000 50000 51000 52000 53000 54000 55000 56000 57000 58000 59000 60000 61000 62000 63000 64000 65000 66000 67000 68000 69000 70000 71000 72000 73000 74000 75000 76000 77000 78000 79000 80000 81000 82000 83000 84000 85000 86000 87000 88000 89000 90000 91000 92000 93000 94000 95000 96000 97000 98000 99000 100000

ABSTRACT OF THE DISCLOSURE

SUPPORT FOR EXHAUSTION RECOVERY IN A DATA PROCESSING SYSTEM WITH MEMORY COMPRESSION

5

A memory exhaustion condition is handled in a data processing system having first and second regions of physical memory. The memory exhaustion condition is detected while the second region is mirroring at least part of the first region. In response to the memory exhaustion condition, memory mirroring is at least partially deactivated and at least part of the second region is utilized to augment the first region, such that the memory exhaustion condition is eliminated. In an illustrative embodiment, the data processing system compresses real memory into the first region of physical memory, and the memory exhaustion condition arises when the first region lacks sufficient available capacity to accommodate current requirements for real memory. The memory exhaustion condition is eliminated by compressing at least part of the real memory into the second region.

10

15

20

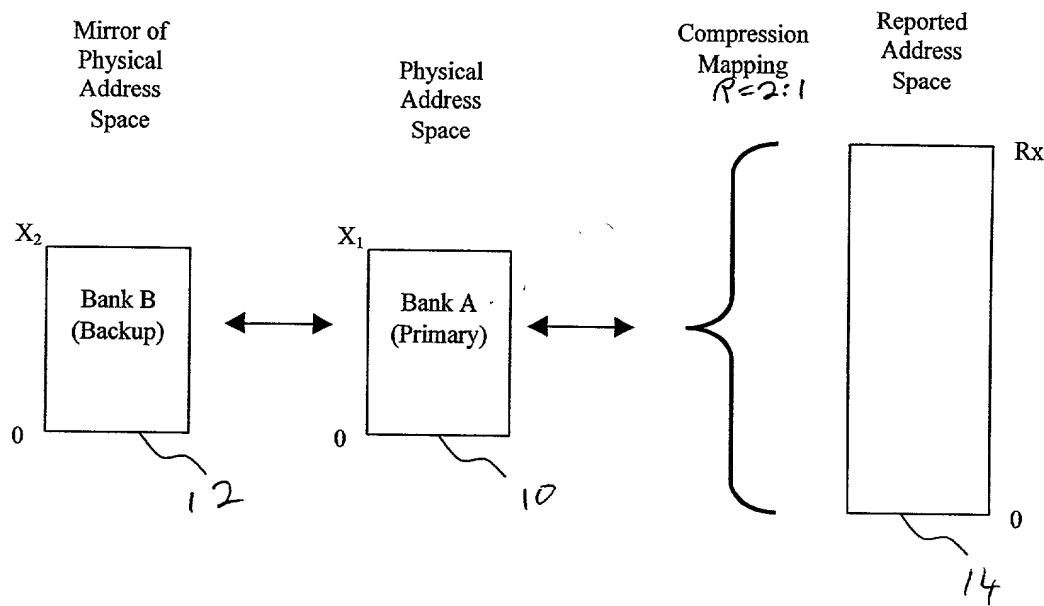


FIG. 1A

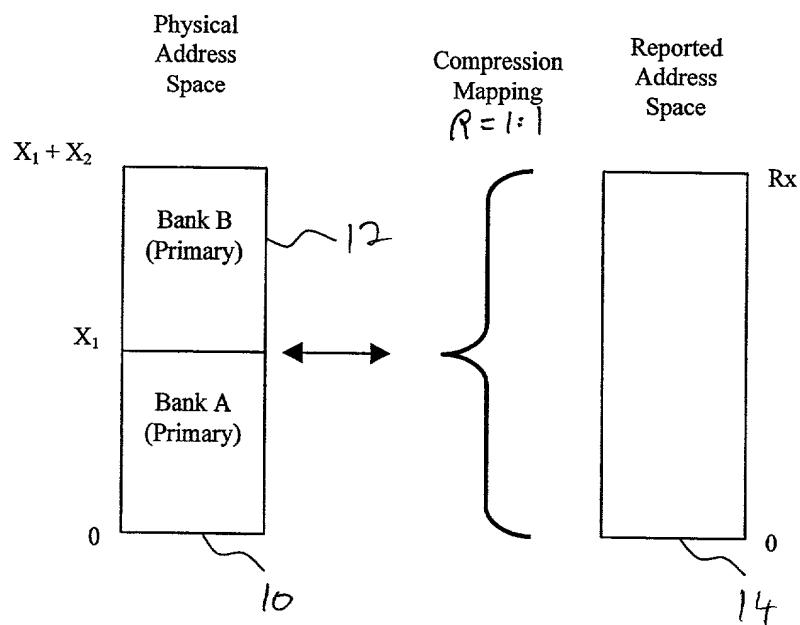


FIG. 1B

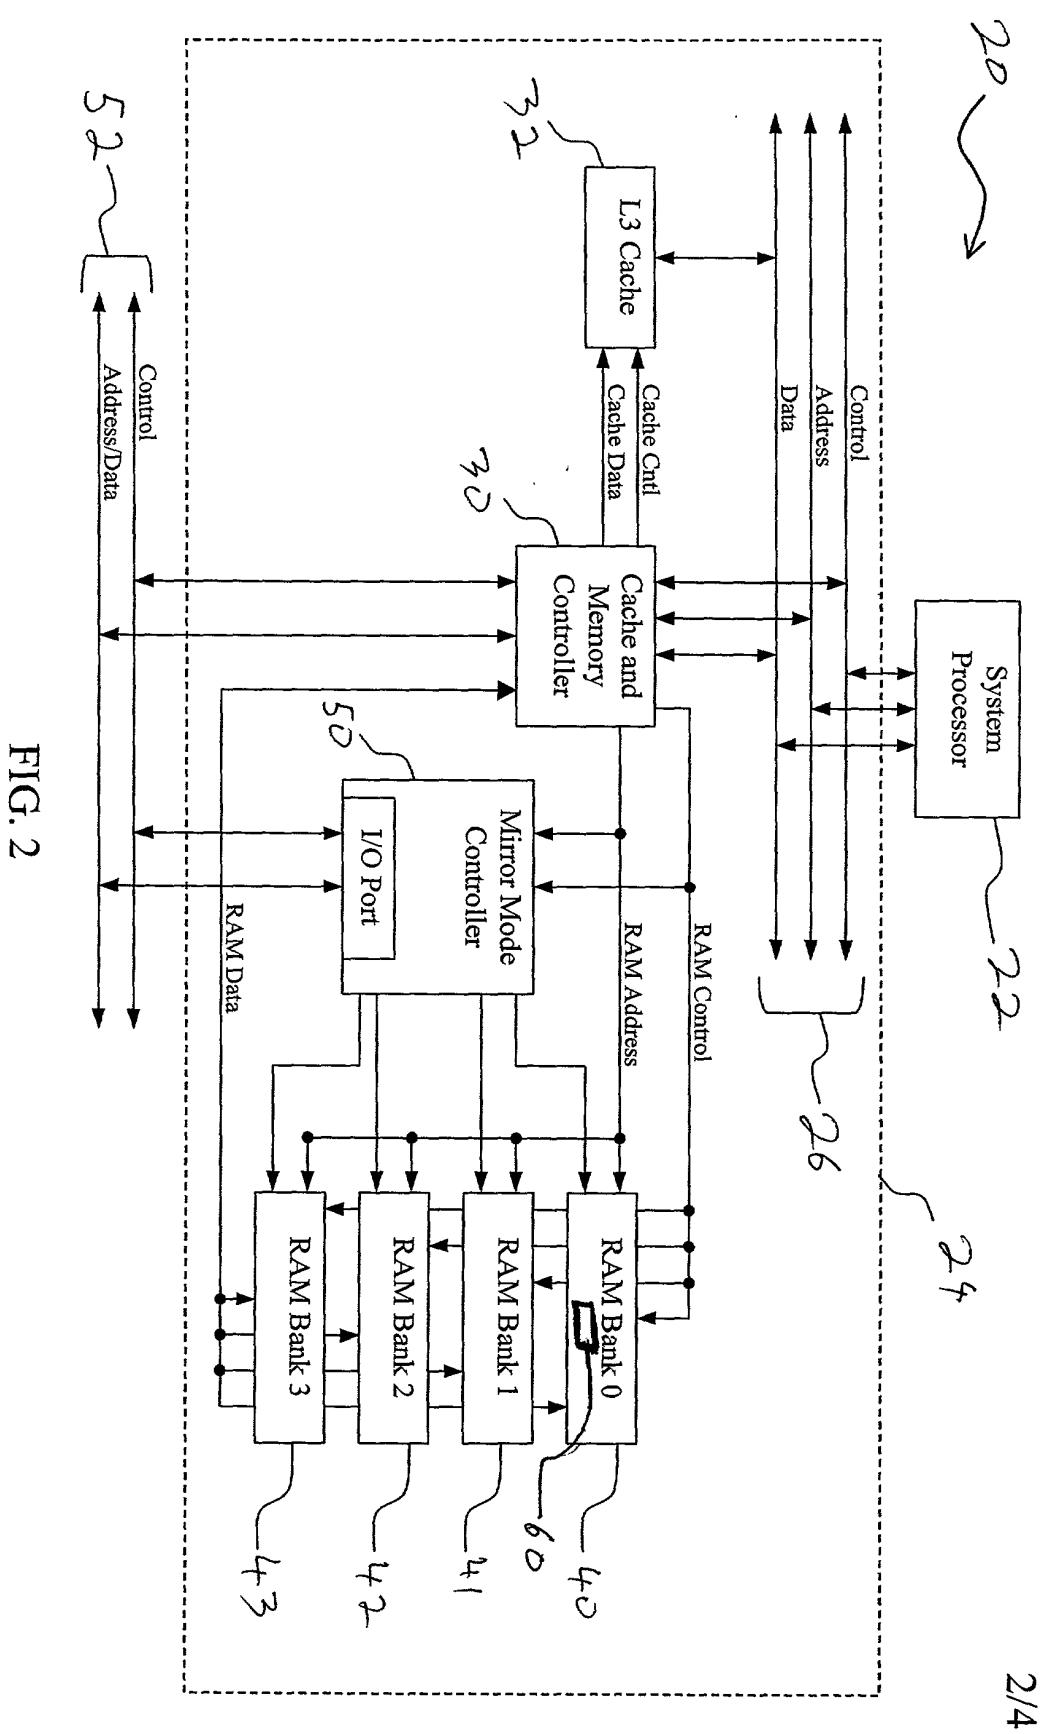


FIG. 2

| BSEL1 | BSEL0 | |
|-------|-------|--|
| 0 | 0 | Mirror Bank0 on Bank3 |
| 0 | 1 | Mirror Bank1 on Bank3 |
| 1 | 0 | Mirror Bank2 on Bank3 |
| 1 | 1 | Use Bank3 as the high bank of a 4-bank array |

FIG. 3

| MR | |
|----|---|
| 0 | Read from primary bank |
| 1 | Read from Bank3 in lieu of primary bank |

FIG. 4

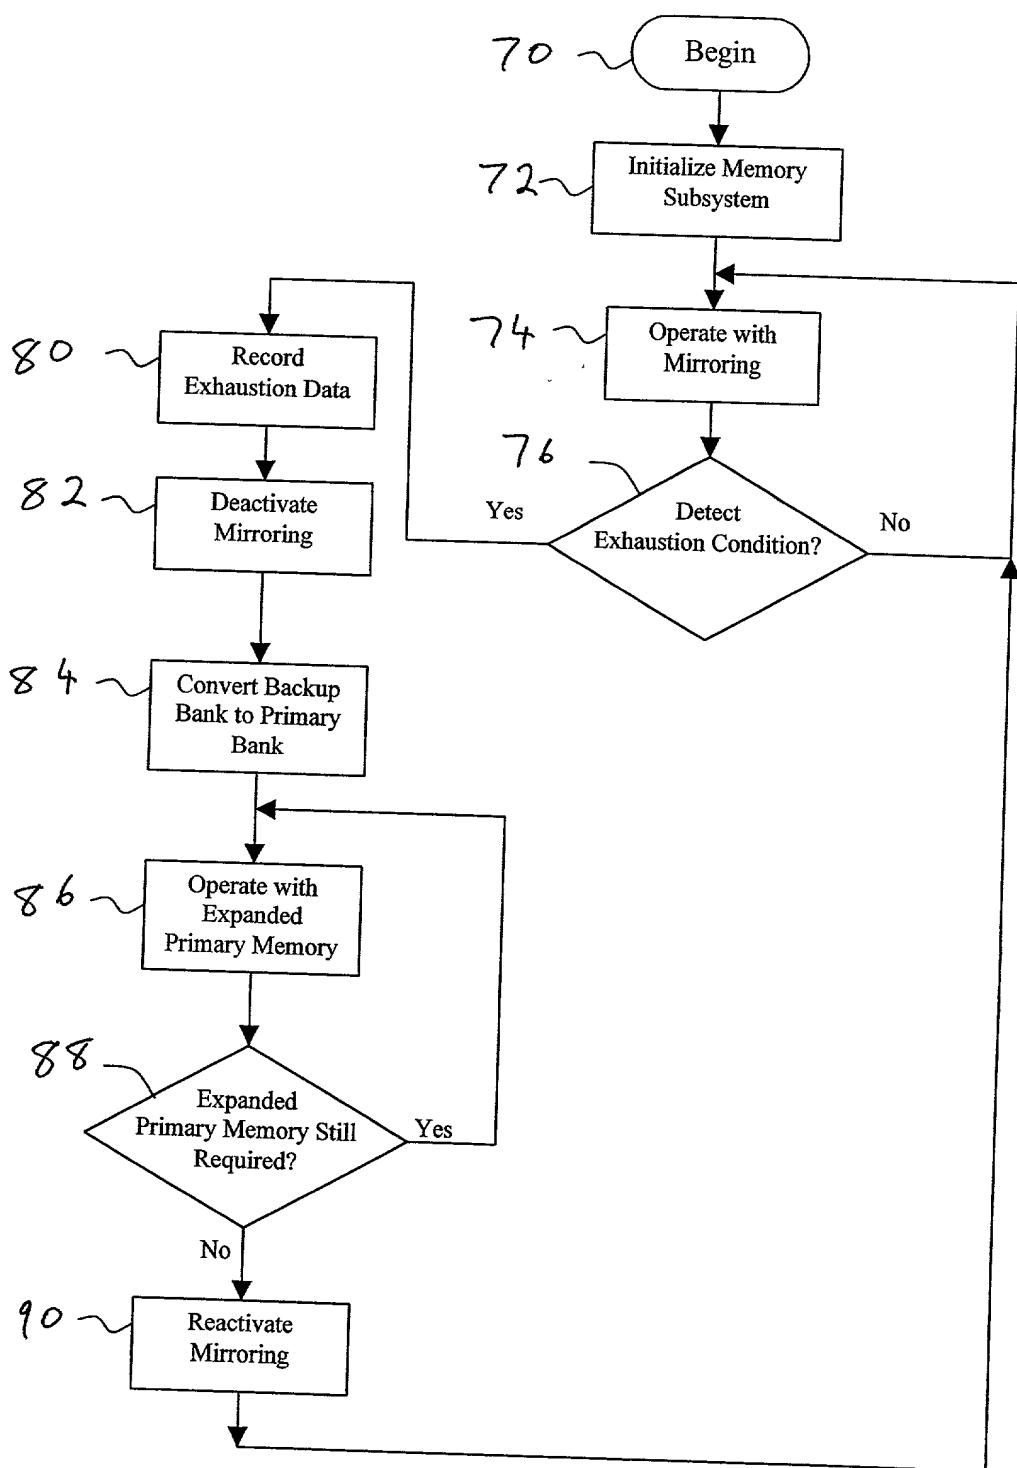


FIG. 5

DECLARATION AND POWER OF ATTORNEY FOR
PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

SUPPORT FOR EXHAUSTION RECOVERY IN A DATA PROCESSING SYSTEM WITH MEMORY MIRRORING

the specification of which (check one)

X is attached hereto.

— was filed on _____
as Application Serial No. _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

| Prior Foreign Application(s): | Priority Claimed | |
|-------------------------------|--------------------|------------------|
| _____ | _____ Yes _____ No | |
| (Number) | (Country) | (Day/Month/Year) |

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal

DOCKET NUMBER: RPS920000076US1

Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

| (Application Serial #) | (Filing Date) | (Status) |
|------------------------|---------------|----------|
|------------------------|---------------|----------|

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Horace St. Julian, Reg. No. 30,329; J. Bruce Schelkopf, Reg. No. 43,901; John D. Flynn, Reg. No. 35,137; George E. Grosser, Reg. No. 25,629; Martin J. McKinley, Reg. No. 31,782; Christopher A. Hughes, Reg. No. 26,914; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Andrew J. Dillon, Reg. No. 29,634; Daniel E. Venglarik, Reg. No. 39,409; Jack V. Musgrove, Reg. No. 31,986; Brian F. Russell, Reg. No. 40,796; Matthew W. Baca, Reg. No. 42,277; Antony P. Ng, Reg. No. 43,427; Michael R. Barre, Reg. No. 44,023; Andrew Mitchell Harris, Reg. No. 42,638; Richard McCain, Reg. No. 43,785; and Michael Noe, Jr., Reg. No. 44,975.

Send correspondence to: Andrew J. Dillon, FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP, Suite 350 Lakewood on the Park, 7600B North Capital of Texas Highway, Austin, Texas 78731, and direct all telephone calls to Andrew J. Dillon, (512) 343-6116.

FULL NAME OF SOLE OR FIRST INVENTOR: Charles David Bauman

INVENTORS SIGNATURE: Charles David Bauman DATE: 10/13/2000

RESIDENCE: 22925 NE 12th Place
Sammamish, Washington 98074

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 22925 NE 12th Place
Sammamish, Washington 98074

FULL NAME OF SECOND INVENTOR: Richard Bealkowski

INVENTORS SIGNATURE: Richard Bealkowski DATE: 10/13/2000

RESIDENCE: 8336 158th Place NE
Redmond, Washington 98052

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 8336 158th Place NE
Redmond, Washington 98052

FULL NAME OF THIRD INVENTOR: Thomas J. Clement

INVENTORS SIGNATURE: Thomas J. Clement DATE: 10/17/2000

RESIDENCE: 12700 Victoria Woods Drive
Raleigh, North Carolina 27613

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 12700 Victoria Woods Drive
Raleigh, North Carolina 27613

FULL NAME OF FOURTH INVENTOR: Jerry William Pearce

INVENTORS SIGNATURE: Jerry William Pearce DATE: 10/17/2000

RESIDENCE: 825 Green Passage Lane
Apex, North Carolina 27502

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 825 Green Passage Lane
Apex, North Carolina 27502

FULL NAME OF FIFTH INVENTOR: Michael Robert Turner

INVENTORS SIGNATURE: Michael Robert Turner DATE: 10/13/2000

RESIDENCE: 29904 NE 52nd Street
Carnation, Washington 98014

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 29904 NE 52nd Street
Carnation, Washington 98014